

REMARKS/ARGUMENTS

The arguments presented herein include the arguments and amendments Applicants discussed with the Examiner during phone interview dated November 20, 2008. The Examiner requested Applicants to submit the discussed arguments for reconsideration, which Applicants present herein. Applicants submit that the arguments presented herein make the substance of the phone interview of record to comply with 37 CFR 1.133. If the Examiner believes that further information on the interview needs to be made of record to comply with the requirements, Applicants request the Examiner to identify such further information.

Applicants amended the Title and Abstract to conform to the method type of claim pending in the Application.

1. Claims 1, 3-5, 8-12, 17-18 are Patentable Over the Cited Art

The Examiner rejected claims 1, 3-5, 8-12, and 17-18 as obvious over Guha (U.S. Pub. No. 2002/0194324) and in view of Ng (U.S. Patent Pub. No. 2004/0049565) and Bradley (U.S. Patent No. 7,082,463). Applicants traverse with respect to the amended claims.

Claim 1 recites a method for managing a network providing Input/Output (I/O) paths between a plurality of host systems and storage volumes in storage systems, comprising: providing an application service connection definition for each of the I/O paths from a host to a storage volume; providing at least one service level guarantee definition indicating performance criteria to satisfy service requirements included in at least one service level agreement with at least one customer for network resources; associating each service level guarantee definition with at least one application service connection definition; gathering, by a virtualization controller mapping physical storage resources to virtual volumes in a virtualization layer, Input/Output (I/O) performance data for I/O requests transmitted through the I/O paths; transmitting, by the virtualization controller, the gathered performance data to a service level agreement server; determining, by the service level agreement server, performance data maintained for the application service connection for which the gathered performance data was received; updating, by the service level agreement server, the determined performance data with the performance data received from the virtualization controller; monitoring, by the service level agreement server, whether the performance data for application service connections indicating the I/O requests transmitted through the I/O paths satisfy the performance criteria indicated in the

service level guarantee definition associated with the application service connection definitions for the I/O path; and transmitting, by the service level agreement server, commands to the virtualization controller to throttle I/O transmission over at least one connection in response to determining that the performance data for at least one application service connection for the connection does not satisfy the performance criteria.

The Examiner cited paras. 60 and 69 of Guha as teaching the claim requirement of transmitting, by the virtualization controller, the gathered performance data to a service level agreement server. The Examiner further modified Guha with the virtualization controller of Ng, citing paras. 2, 23, 34, and 43 of Ng because the Examiner recognized that Guha does not disclose a virtualization controller. (OA4, pg. 3) Applicants traverse.

The cited para. 60 of Guha discusses how a content controller uses content request information from the quality of service (QoS) enforcer captured at the entry point of data center to allocate or deallocate content storage. The cited para. 69 discusses how the QoS enforcer communicates with the content controller. The content controller maintains and controls metadata associated with content data. The combination of the QoS enforcer and content controller allow dynamic allocation of I/O resources based on I/O load. If the traffic increases, the content controller might create and allow access to replicate web content.

Although the cited Guha discusses how a content controller uses content request information from a QoS enforcer, this does not teach transmitting gathered performance data. Instead, the cited Guha discusses that the QoS enforcer captures content request information. Further, there is no teaching that a virtualization controller, which maps physical storage to virtual volumes, gathers I/O performance data and transmits to a service level agreement (SLA) server.

The cited paras. 2 and 34 of Ng mentions that storage devices can share their storage resources using virtualization controllers that make disks appear as one or more logical units. The para. 23 of Ng mentions that aspects of measuring data access and data flow over a fiber channel switch may be designed into the firmware of a virtual storage exchange device. Para. 34 mentions that Para. 43 mentions that the virtual devices (VSX) collect information used for flow control, load balancing, operating the SAN, and to provide a required quality of service (QoS) under service level agreements.

Although the cited Ng discusses a virtualization device collecting information for flow control and for quality of service purposes, the Examiner has not cited any part of Ng or Guha that teaches or suggests that a virtualization controller transmits gathered performance data to a service level agreement server.

The Examiner cited para. 69 of Guha, modified by Ng, as teaching the claim requirement the service level agreement server transmitting commands to the virtualization controller to throttle I/O transmission if the performance data for one connection does not satisfy performance criteria. (OA4, pg. 3) Applicants traverse.

The cited para. 69 discusses how the QoS enforcer communicates with the content controller. The content controller maintains and controls metadata associated with content data. The combination of the QoS enforcer and content controller allow dynamic allocation of I/O resources based on I/O load. If the traffic increases, the content controller might create and allow access to replicate web content.

The Examiner has not explained which element of the cited Guha should be modified to be a virtualization controller as mentioned in Ng, either the QoS enforcer or the content controller. Moreover, the Examiner has not shown how if either the QoS enforcer or the content controller of Guha is modified by Ng to be a virtualization controller, that the QoS enforcer or content controller so modified as proposed teaches the claim requirements that a virtualization controller (i.e., one of the cited QoS enforcer or content controller) transmits performance data to a service level agreement server (i.e., one of the cited QoS enforcer or content controller), such that the service level agreement server transmits commands to the virtualization controller to throttle I/O transmission over a connection if the transmitted performance data does not satisfy the performance criteria.

In other words, the Examiner has not shown where Guha teaches that the either the QoS enforcer or content controller transmits performance data to the other, so the other component can transmit commands to cause the QoS enforcer or content controller to throttle I/O transmission over a connection if the transmitted performance data does not satisfy performance criteria. Instead, the combination yields a system where a QoS enforcer performs load balancing and makes routing decisions and communicates with a content controller that uses information from the QoS enforcer to allocate content storage (Guha) and where a virtualization controller may also collect information for quality of service and flow controller. However, the

cited Ng and Guha do not teach or suggest that the virtualization controller and service level agreement interact as claimed to maintain a quality of service for different application service connections.

Thus, even if one modifies the QoS enforcer or content controller of Guha with the virtualization controller of Ng, the claim requirements with respect to a virtualization controller and SLA server are still not taught or suggested by the cited combination of Guha and Ng.

The Examiner found that the combination of Guha and Ng do not disclose the claim requirement of a service level agreement server determining performance data maintained for the application service connection for which the gathered performance data was received; updating, by the service level agreement server, the determined performance data with the performance data received from the virtualization controller. The Examiner cited col. 2, line 60 to col. 3, line 8 of Bradley as teaching this claim requirement. (OA4, pg. 4) Applicants traverse.

The cited cols. 2-3 of Bradley mentions a time based service monitoring mechanism for monitoring service level agreements to monitor a level of network service offered by a service provider. Data is received for defining one or more tests for monitoring the level of network service provided to a customer. The tests are distributed to agents configured to communicate with devices.

Although the cited Bradley discusses how service level agreements are monitored, there is no teaching in Bradley or the cited Guha-Ng, that a service level agreement server updates performance data with performance data received from a virtualization controller, and then use that performance data to determine whether performance criteria is satisfied to determine whether to send commands to the virtualization controller that provided the performance data to have the virtualization controller throttle I/O transmission for a connection for which performance criteria is not satisfied.

Accordingly, Applicants submit that amended claim 1 is patentable over the cited combination because the cited combination of Guha and Ng do not teach or suggest all the requirements of claim 1.

Claims 2-5, 8-12, 17, and 18 are patentable over the cited combination of Guha, Ng, and Bradley because they depend from claim 1, which is patentable over the cited art for the reasons discussed above, and the additional requirements of these claims in combination with the base

claims provide further grounds of patentability over the cited art. The following claims provide additional grounds of patentability over the cited art for the below discussed reasons.

Claim 2 depends from claim 1 and further requires that each service level guarantee definition is implemented as a separate element in at least one Extended Markup Language (XML) document, the element for the service level guarantee definition includes the performance criteria defined in the service level agreement, and wherein the application service connection definition for each of the I/O paths is implemented as an element the at least one XML document, wherein attributes of the application service connection definition element provide information on the I/O path.

The Examiner cited col. 12, Table 4, col. 9, line 29, col. 7, lines 5-17, col. 13, table 5 of Bradley as teaching the requirements of claim 2. (OA4, pgs. 4-5) Applicants traverse.

The cited Table 4 shows a service level agreement component of a type of metric to be monitored. The cited col. 9 mentions a service level contract. The cited table 5 discusses a metric component used to measure network latency between two devices. See, Bradley, col. 12, lines 60-66.

Although the cited Bradley discusses service level agreements and information for such agreements and metrics, there is no teaching that each service level agreement is implemented as a separate element in an XML document and that an application service connection for each I/O paths is also implemented as an element. The cited Bradley does not teach or suggest the specific requirements for implementing an SLA and application service connection information as elements in an XML document.

Accordingly, claim 2 provides additional grounds of patentability over the cited art because the additional requirements of claim 2 are not taught or suggested in the cited references.

Claim 3 depends from claim 1 and further requires that multiple service level guarantee definitions indicating different performance criteria are associated with different sets of application service connection definitions.

The Examiner cited para. 49 of Guha as teaching the additional requirements of claim 3. (OA4, pg. 5) Applicants traverse.

The cited para. 49 mentions that a rules table 32 contains a QoS policy for each address, such as a priority, response time, etc. A resource status contains information from content controller 36 indicating the status of the components.

Applicants submit maintaining a QoS policy for each address does not teach the claim requirement that multiple service level guarantee definitions indicating different performance criteria are associated with different sets of application service connection definitions.

Accordingly, claim 3 provides additional grounds of patentability over the cited art because the additional requirements of claim 3 are not taught or suggested in the cited references.

Claim 4 depends from claim 3 and further requires that the application service definition for the I/O paths may be associated with the multiple service level guarantee definitions, wherein the monitoring comprises determining whether the I/O requests transmitted through the I/O paths satisfy the performance criteria of all associated service level guarantee definitions.

The Examiner cited the above discussed para. 49 of Guha as teaching the additional requirements of claim 4. (OA4, pg. 5) Applicants traverse.

The cited para. 49 mentions rules table 32 contains a QoS policy for each address, such as priority, response time, etc. A resource status contains information from content controller 36 indicating the status of the components.

Nowhere does the cited para. 49 teach that multiple SLAs are associated with an application service definition for I/O paths and that the monitoring determines whether all performance criteria of all SLAs for an application service definition are satisfied. Instead, the cited para. 49 mentions how a QoS policy is provided for each address.

Accordingly, claim 4 provides additional grounds of patentability over the cited art because the additional requirements of claim 4 are not taught or suggested in the cited references.

Claim 5 depends from claim 1 and further requires providing an application service group identifying a plurality of application service connection definitions, wherein associating the at least one service level guarantee definition with the application service connection definitions comprises associating the at least one service level guarantee definition with the at least one application service group, wherein the application service connection definitions identified in the application service group are associated with the service level guarantee definition with which their application service group is associated.

The Examiner cited para. 14 of Guha as teaching the additional requirements of claim 5. (OA4, pgs. 5-6) Applicants traverse.

The cited para. 14 mentions that a QoS enforcer has a rule engine containing a predetermined QoS policy, and a global infrastructure control (GIC) determines which data

storage centers provide data to meet a content request according to the QoS policy and status information.

Although the cited para. 14 discusses how a GIC determines whether data storage centers meet a QoS policy, there is no teaching of associating a SLA definition with application service groups that further identify application service connections. Instead, the cited para. 14 discusses determining whether storage centers satisfy a QoS policy.

Accordingly, claim 5 provides additional grounds of patentability over the cited art because the additional requirements of claim 5 are not taught or suggested in the cited references.

Claim 8 depends from claim 1 and further requires that the monitoring whether the I/O requests transmitted through the I/O path satisfy performance criteria indicated in the service level guarantee definition comprises: gathering performance information concerning I/O requests for the I/O paths; selecting one of the at least one service level guarantee definition; and for each of the I/O paths identified by the application service connection definition associated with the selected service level guarantee definition, comparing the gathered performance information for the I/O path with the performance criteria indicated in the selected service level guarantee definition.

The Examiner cited paras. 44, 45, 60, and 62 of Guha as teaching the additional requirements of these claims. (OA4, pg. 6) Applicants traverse.

The cited para. 44 mentions that content requests are monitored by a Quality of Service (QoS) enforcer that tracks every content request to a file server. The cited para. 45 discusses a Quality of Service (QoS) enforcer that makes routing decisions to provide load balancing, based on a rule based system that associates a QoS policy with content. Information on application servers and their loads are provided to the QoS enforcer. The cited paras. 60 and 62 discuss how a content controller uses content request information from the QoS enforcer captured at the entry point of data center to allocate or deallocate content storage.

Although the cited Guha discusses tracking content requests and a QoS enforcer making routing decisions to provide load balancing based on a rules based system and QoS policy, there is no teaching of comparing, for each I/O path in an application service connection definition, performance information for the I/O path with the performance criteria indicated in the SLA definition. Although Guha discusses determining whether performance information for I/O paths in an application service

connection definition satisfy performance criteria in the service level guarantee definition associated with the application service connection definition. Instead, the cited Guha discusses using content request information to determine to allocate or deallocate content storage.

Accordingly, claim 8 provides additional grounds of patentability over the cited art because the additional requirements of claim 8 are not taught or suggested in the cited references.

Claims 9-15 depend directly or indirectly from claim 8 and provide further requirements concerning how operations are adjusted among I/O paths and how performance information for I/O paths is gathered concerning I/O response time and I/O demand. These claims provide further details, which in combination with base and intervening claims, are also not taught or suggested in the cited Guha and other references.

Conclusion

For all the above reasons, Applicant submits that the pending claims 1-15, 17, and 18 are patentable. Should any additional fees be required beyond those paid, please charge Deposit Account No. 09-0466.

The attorney of record invites the Examiner to contact him at (310) 553-7977 if the Examiner believes such contact would advance the prosecution of the case.

Dated: November 28, 2008

By: /David Victor/

David W. Victor
Registration No. 39,867

Please direct all correspondences to:

David W. Victor
Konrad Raynes & Victor, LLP
315 South Beverly Drive, Ste. 210
Beverly Hills, CA 90212
Tel: (310) 553-7977
Fax: 310-556-7984